

REMARKS

Claims 1-21 are pending in the present application. No amendments to the claims have been made by this Response. Reconsideration of the claims is respectfully requested.

Amendments were made to the specification to update information for the cross-referenced applications. No new matter has been added by any of the amendments to the specification.

I. Telephone Interview

Applicants thank Examiner Patel and his supervisor for the courtesies extended to Applicants' representative during the June 7, 2004 telephone interview. During the interview, Applicants' representative discussed the distinctions of the present claims over the cited art. Specifically, Applicants' representative asserted that neither Krause et al. nor Shah, et al., either alone or in combination, teach or suggest storing a reliable datagram domain within a reliable datagram queue pair context, storing the same reliable datagram domain within an end-to-end context, or storing a partition key within the end-to-end context. In response, the Examiners stated that they would not read limitations from the specification into the claims with regard to the definitions of these terms specifically provided by the specification.

Applicants' representative asserted that the Examiner must interpret the claims in light of the specification in order that the Examiner may determine whether the specific features in the claims, or their equivalents, are taught or suggested in the prior art. Applicants are not requesting that features be read into the claims but rather, that the Examiner use the definition of the terms already present in the claims. The definitions of these terms are intrinsic in the use of the terms in the claims. No new features are added to the claim by merely interpreting the claim terms in the way they are defined in the specification.

The only way that the Examiner can determine whether a feature or its equivalent is taught in the prior art is to first know what the feature is. In the present case, this

means going to the specification to determine what a reliable datagram domain is, what a reliable datagram queue pair context is, and what an end-to-end context is. This means that the Examiner must read these terms as having the meaning set forth by the definitions in the specification. This does not introduce new features into the claims. Furthermore, even if the Examiners do not go to the specification to determine what is meant by these terms, the prior art still does not teach or suggest these features as detailed below. The substance of the telephone interview is summarized in the following remarks.

II. Objection to Specification

The Office Action objects to the specification stating that the related applications on page 1 of the application are missing their corresponding serial numbers. By this Response, the specification is amended to update the cross-referenced application information. Accordingly, Applicants respectfully request withdrawal of the objection to the specification.

III. 35 U.S.C. § 103, Alleged Obviousness Based on Krause and Shah

The Office Action rejects claims 1, 8-10, and 17-19 under 35 U.S.C. § 103(a) as being allegedly unpatentable over Krause et al. (International Publication No. WO 00/72159) in view of Shah et al. (U.S. Patent No. 6,694,361). This rejection is respectfully traversed.

As to claim 1, the Office Action states:

Referring to claim 1,

The reference Krause teaches the reliable datagram services with an underlying end-to-end context of a channel adapter. (Fig.6) The reference also teaches that the SDRs at the source and SDRs at the destination are created where specific application instances associated with specific SDRs. (storing Reliable datagram domain within reliable datagram and end-to-end context)(page 3, lines 5-15, Fig.4). The reference teaches the examples of protection violations including invalid protection keys, invalid access right request etc. (page 21, lines 30-31 and page 22 lines 1-4). The reference fails to teach the queue pairs and storing partition keys

within end-to-end contexts. The reference Shah teaches the virtual interface which can contain work queues formed in pairs including a send queue and a receive queue which is located in channel adapters of any host systems. (col.3, lines 29-47, col.6, lines 22-43). The reference also teaches the subnet manager responsible for various functions at the channel adapter including assigning unique addresses to all channel adaptor ports. Also, the reference teaches partition manager as part of subnet manager assigning storing partition keys to the fabric agent (channel adapter ports)(col.7, lines 23-42) Therefore, it would have been obvious for one in ordinary skill in the art at the time the invention was made to modify Krause to include the teachings of Shah such that the each SDR is assured a reliable connection with the incoming datagram by a partition key assigned to the fabric agent (channel adapter ports). This will allow the re-use of the same destination SDR resource set per multiple destination application instances particularly identified by the partition keys in the datagram.

Office Action dated March 8, 2004, pages 2-3.

Claim 1, which is representative of the other rejected independent claims 10 and 19 with regard to similarly recited subject matter, reads as follows:

1. A method for associating reliable datagram queue pairs with an underlying end-to-end context of a channel adapter, comprising:
storing a reliable datagram domain within the reliable datagram queue pair context;
storing the same reliable datagram domain within the end-to-end context; and
storing a partition key within the end-to-end context.

(emphasis added)

Neither Krause nor Shah, either alone or in combination, teach or suggest a reliable datagram domain, a reliable datagram queue pair context, an end-to-end context, or a partition key, let alone the specific features of storing a reliable datagram domain within a reliable datagram queue pair context, storing the same reliable datagram domain within an end-to-end context, or storing a partition key within an end-to-end context.

Krause is directed to are reliable datagram service in which source and destination resources (SDRs) having SDR resources are used to multiplex units of work produced by at least one source application instance (AI) into a serial unit of work stream having units of work in a defined order. The serial unit of work stream is then transmitted over a

communication services fabric. The SDR resources provide the communication protocol for acknowledging and negatively acknowledging units of work in the serial unit of work stream. The same SDR resource set may be used for multiple destination AIs. In addition, any number of source AIs can employ an established SDR to communicate to any number of destination AIs.

While Krause provides a reliable datagram service, there is nothing in Krause that teaches a reliable datagram domain, a reliable datagram queue pair context, an end-to-end context, or a partition key, as recited in claim 1. It is well established that the specification serves as a dictionary for the terms used in the claims. Thus, the claims must be interpreted in light of the specification. As defined in the present specification, an end-to-end context (EEC) is a host channel adapter (HCA) facility for providing reliability attributes for communication between at least two nodes (page 23, line 29 to page 24, line 3). The reliable datagram domain (RDD) is a data structure field that associates a reliable datagram queue pair with a particular EEC and provides a means for restricting access to a specific EEC (page 5, lines 7-10). A queue pair context is one or more blocks of information that contain pointers to, length, and other information needed to manage the individual queues in a queue pair (page 16, lines 18-20). Krause does not teach or suggest any such structures. Furthermore, Krause does not teach or suggest the specific features of claim 1 which operate on such structures, i.e. Krause does not teach or suggest storing a reliable datagram domain within a reliable datagram queue pair context, storing the same reliable datagram domain within an end-to-end context, or storing a partition key within an end-to-end context.

The Office Action alleges that Krause teaches storing a reliable datagram domain within reliable datagram and end-to-end contexts at page 3, lines 5-15 and in Figure 4. Page 3, lines 5-15 read as follows:

The present invention provides a source and destination resource (SDR) having source SDR resources at a source device and destination SDR resources at a destination device. The source SDR resources multiplex units of work produced by at least one source application instance (AI) into a serial unit of work stream having units of work in a defined order and transmit the serial unit of work stream over a communication services/fabric. The destination SDR resources receive

the serial unit of work stream, demultiplex the serial unit of work stream into units of work provided to at least one destination AI, and provide a negative acknowledgement (NAK) for a unit of work received ahead of its defined order.

Nowhere in this section, or any other section, of Krause is there any teaching or suggestion of a reliable datagram domain, a reliable datagram context, or an end-to-end context, let alone the specific features of claim 1 that make use of these structures. All that this section of Krause teaches is that SDR resources may be used to multiplex units of work from a plurality of source application instances into a serial unit of work stream and may be used to demultiplex the serial unit of work stream into units of work provided to at least one destination AI. There is no teaching anywhere in Krause regarding a reliable datagram domain, a reliable datagram context, or an end-to-end context, despite the allegations made by the Office Action to the contrary.

Furthermore Figure 4 of Krause does not teach or suggest any of these structures either. Figure 4 of Krause merely illustrates how a plurality application instances on a plurality of devices may communicate over a communication services/fabric using source and destination resources (SDRs). Nowhere in Figure 4 is there any reliable datagram domain, reliable datagram context, or end-to-end context depicted. This is because there is no teaching in the actual text of Krause regarding any of these structures, let alone the specific features of storing a reliable datagram domain within a reliable datagram queue pair context, storing the same reliable datagram domain within an end-to-end context, or storing a partition key within an end-to-end context.

Shah, likewise, does not teach or suggest these features. Shah teaches a methodology for assigning multiple local identifiers (LIDs) to ports in a cluster. With the method of Shah, a subnet manager performs a topology discovery of the cluster and detects ports associated with the fabric. The subnet manager then computes a minimal spanning tree for the cluster which connects every port to every other port through a single path. The subnet manager then assigns a single base LID to each port and programs the assigned LIDs into forwarding tables in the fabric. Additional LIDs are reserved for each port while the single base LID is assigned to each port such that the

baric is functional and connected when the subnet manager is performing a path analysis of the cluster.

Shah teaches work queues formed in pairs including send queues and receive queues which are located in channel adapters of a host system (column 3, lines 29-47 and column 6, lines 22-43). Shah also teaches a subnet manager that assigns unique addresses to all channel adapter ports. Shah further teaches a partition manager as part of the subnet manager that assigns partition keys to the fabric agents (column 7, lines 23-42). However, nowhere in Shah is there any teaching or suggestion regarding storing a reliable datagram domain within a reliable datagram queue pair context, storing the same reliable datagram domain within an end-to-end context, or storing a partition key within an end-to-end context. While Shah may teach channel adapters, work queue pairs, subnet managers, LIDs, and a plethora of other system area network (SAN) components, there is absolutely no mention what-so-ever of any of the features recited in claim 1.

Thus, neither Krause nor Shah, whether taken alone or in combination, teach or suggest the specific features recited in claim 1. Therefore, any alleged combination of Krause and Shah, assuming that such a combination were possible and that one of ordinary skill in the art would somehow be motivated to make such a combination, still would not result in the invention as recited in claim 1.

Moreover, one of ordinary skill in the art, in actuality, would not have been motivated to combine Shah with Krause in the manner alleged by the Office Action. Shah is directed to a system for assigning multiple local identifiers to ports in a subnet or cluster of a system area network. Krause is directed to a mechanism for communicating between a plurality of application instances using source destination resources. There is no problem stated in either reference for which the other reference is a solution. In other words, neither reference teaches or suggests the desirability of including the teachings of the other reference.

The Office Action alleges that the reason one of ordinary skill in the art would be motivated to combine Shah with Krause is so that "each SDR is assured a reliable connection with the incoming datagram by a partition key assigned to the fabric agent (channel adapter ports). This will allow the re-use of the same destination SDR resource set per multiple destination application instances particularly identified by the partition

keys in the datagram." There is no support for this allegation in the actual references. That is, there is nothing in Krause that teaches any need or desirability to include partition keys assigned to fabric agents. Moreover, there is not teaching in Shah to include the ability to reuse the "same destination SDR resource set per multiple destination application instances" or even how such features might be used in conjunction with one another. The alleged motivation is simply a hindsight attempt to validate a combination of references that actually are either not combinable or would not be combined by those of ordinary skill in the art. That is, the only way that one of ordinary skill in the art would be motivated to combine the references in the manner alleged by the Office Action, and to make the necessary modifications to the references to arrive at the claimed invention, is if that person of ordinary skill in the art had a prior knowledge of Applicants' claimed invention and the sole intent to recreate the claimed invention. This is clearly a hindsight reconstruction that is not a permissible basis upon which to make a rejection under 35 U.S.C. § 103(a).

In view of the above, Applicants respectfully submit that neither Krause nor Shah, either alone or in combination, teach or suggest the features of independent claims 1, 10 and 19. At least by virtue of their dependency on claims 1, 10 and 19, respectively, Krause and Shah also do not teach or suggest the features of dependent claims 8-9, and 17-18. Accordingly, Applicants respectfully request withdrawal of the rejection of claims 1, 8-10, and 17-19 under 35 U.S.C. § 103(a).

In addition to the above, neither Krause nor Shah, either alone or in combination, teach or suggest the specific features recited in dependent claims 8-9 and 17-18. For example, with regard to claims 8 and 17, neither reference teaches or suggests storing reliable datagram domain numbers for kernel code and user code, wherein the kernel reliable datagram domain can only be used by kernel code.

The Office Action admits that Krause does not teach this feature but alleges that this feature is taught by Shah simply because Shah teaches bypassing the kernel and accessing a host channel adapter directly by the users or the consumer processes (column 5, lines 40-67 or Shah). Simply teaching that the kernel can be bypassed does not teach or suggest storing reliable datagram domain numbers for kernel code and user code or that the kernel reliable datagram domain can only be used by kernel code. All that Shah

teaches is that there is a kernel and there are user or consumer processes and that the user or consumer processes may bypass the kernel to access the host channel adapter. There is nothing in Shah that teaches to use a kernel code reliable datagram domain number or user code reliable datagram domain number to do this bypassing of the kernel. There is also nothing in Shah that teaches or suggests that a kernel code reliable datagram domain can only be used by kernel code. The Office Action simply has not found these features in any of the cited references.

Moreover, the Office Action's statement of how the references are to be modified and the reason such modifications would be obvious does not even address the specific features of claims 8 and 17. To the contrary, the Office Action states that it would have been obvious to modify Krause to include the teachings of Shah "such that the each SDR (Reliable datagram domain) is directly accessible by the user or the consumer processes by bypassing the operating system kernel process (kernel reliable datagram domain can only be used by kernel code)." This in no way addresses why it would be obvious to modify Krause and Shah such that a reliable datagram domain number is stored for kernel code and for user code. Moreover, despite the Office Action's parenthetical statement, simply bypassing the kernel does not in itself teach or suggest that a kernel reliable datagram domain can only be used by kernel code.

Regarding claims 9 and 18, neither reference, either alone or in combination, teaches or suggests storing reliable datagram domain numbers for kernel code and consumer processes, wherein the kernel reliable datagram domain can only be used by kernel code. Again, for the same reasons as set forth above with regard to claims 8 and 17, neither of the references teach or suggest these features and the Office Action has failed to address why it would be obvious to modify the references to specifically include these features. The alleged motivation and manner of modification does not address these specific features. Thus, in addition to being dependent upon claims 1 and 10, neither Krause nor Shah, either alone or in combination, teach or suggest the specific features of claims 8-9 and 17-18.

IV. 35 U.S.C. § 103, Alleged Obviousness Based on Krause, Shah and Murayama

The Office Action rejects claims 2-7, 11-16, 20 and 21 under 35 U.S.C. § 103(a) as being allegedly unpatentable over Krause in view of Shah and further in view of Murayama et al. (U.S. Patent No. 5,617,424). This rejection is respectfully traversed.

This rejection is based on the same misinterpretation of Krause and Shah addressed above. As discussed above, neither Krause nor Shah, either alone or in combination, teach or suggest the features of independent claims 1, 10 and 19. Murayama does not provide the necessary teachings to provide for the deficiencies noted above with regard to Krause and Shah. That is, Murayama does not teach or suggest storing a reliable datagram domain within a reliable datagram queue pair context, storing the same reliable datagram domain within an end-to-end context, or storing a partition key within an end-to-end context.

Murayama is directed to a method for communicating between network computers by dividing packet data into parts for transfer to respective regions. With the system of Murayama, received data may be transferred directly to the user data region thereby eliminating the need to perform data copy operations. A port ID identifies the region in which the packet is to be received. A port table is used to determine the region which is to receive the data. Murayama makes not mention of a reliable datagram domain, reliable datagram queue pair context, end-to-end context, or a partition key, let alone the specific features that operate on these structures as recited in the independent claims 1, 10 and 19. Thus, even if Murayama were combined with Krause and Shah, and assuming such a combination were possible and one of ordinary skill in the art were somehow motivated to make the combination, the result still would not be the invention as recited in the independent claims 1, 10 and 19. Accordingly, the alleged combination would not teach or suggest every feature of the dependent claims 2-7, 11-16, 20 and 21.

In addition, despite the allegations made by the Office Action, Murayama, in actuality, does not teach or suggest the specific features recited in dependent claims 2-7, 11-16, 20 and 21. For example, with regard to claims 2 and 11, Murayama does not teach that a consumer process cannot directly access a reliable datagram domain. As

mentioned above, none of the references teach or suggest a reliable datagram domain and thus, they could not possibly teach this feature in claims 2 and 11.

The Office Action alleges that Murayama teaches this feature simply because Murayama allegedly teaches that a key in a received packet is checked against a key in the communication region and thus, "a consumer process cannot directly access a reliable datagram domain without protection being checked by comparing the key." In Murayama, the key is used to determine an access right to a "region." A "region" is a data region, i.e. a region of memory. It is not a reliable datagram domain. Furthermore, the key in Murayama does not cause a consumer process to not directly access a reliable datagram domain. Rather, the key is merely a security mechanism for access a data region, i.e. a region of memory.

With regard to claims 3, 12 and 20, Murayama does not teach or suggest comparing a partition key of an incoming packet with the partition key of an end-to-end context, comparing a reliable datagram domain of a queue pair with a reliable datagram domain of the end-to-end context if the partition keys match, or processing a data packet normally if the reliable datagram domains match. The Office Action alleges that Murayama teaches this feature simply because Murayama teaches processing a data packet continuing the information on a receiving region and a security key for protection. The Office Action further states that the receiving side checks the key in the packet that has arrived and the key in the communication region and after the protection key checking is successful, the packet is transferred to a physical address according to the receive region assignment. Based on this, the Office Action somehow comes to a conclusion that the features of claims 3, 12 and 20 are taught by the combination of Krause, Shah and Murayama. However, in actuality, there is nothing in Murayama that teaches anything regarding a reliable datagram domain, an end-to-end context, or a reliable datagram domain of a queue pair. The alleged teaches of Murayama cited by the Office Action have nothing to do with these features, let alone the specific mechanisms of claims 3, 12 and 20 that operate on these features. The Office Action has failed to show where any of the references teach or suggest these features.

Similar distinctions apply to claims 4-7 and 13-16 and 21. Each of these claims make reference to a reliable datagram domain and an end-to-end context, neither of

which are taught or suggested by any of the cited references. Furthermore, with regard to claim 4, none of the cited references, and especially Murayama, teach or suggest processing a data packet according to InfiniBand partitioning semantics, if the partition keys do not match. Regarding claim 5, none of the cited references teach or suggest giving a negative acknowledgment to the data packet if the reliable datagram domains do not match or placing a send queue in an error state. With regard to claim 6, none of the cited references teach or suggest comparing the reliable datagram domain of the queue pair with a reliable datagram domain of an end-to-end context or inserting the partitioning key of the end-to-end context into the transport header of the data packet if the reliable datagram domains match. Regarding claim 7, none of the cited references teach or suggest completing a packet in error if the reliable datagram domains do not match or placing the send queue in an error state.

With regard to each of these rejections, the Office Action makes reference to equally irrelevant portions of Murayama, Krause and Shah references as discussed above with regard to claims 1, 2, 3, 8-10 and 17-19. There simply is nothing in any of the cited references that teaches or suggests any of the features of any of the claims. Nowhere in any of the references is there any teaching of a reliable datagram domain or an end-to-end context, let alone the specific features of claims 4-7. It is noted that nowhere in the Office Action does the Examiner actually point to any portion of the reference that actually teaches a reliable datagram domain or an end-to-end context. Rather, all the sections cited by the Office Action, and summarized by the Office Action, teach other things other than the specific features of a reliable datagram domain and an end-to-end context. For example, with regard to claim 6, the Office Action alleges that Murayama teaches determination of a key that designates an access right and that an outgoing packet is assigned with assigned receiving region and security key. The Office Action then miraculously concludes that Murayama teaches comparing a reliable datagram domain of a queue pair with a reliable datagram domain of an end-to-end context and inserting a partition key of an end-to-end context into a transport header of a data packet if the reliable datagram domains match. There is no way that merely using a key to determine access rights to a memory region or assigning a key to a data packet teaches the specific features recited in claim 6. This is exemplary of the erroneous leaps in logic that

permeate every rejection in the Office Action. Nowhere in any of the references is there any teaching or suggestion that remotely resembles the features of claims 4-7.


Thus, in view of the above, Applicants respectfully submit that neither Krause, Shah nor Murayama, either alone or in combination, teach or suggest the features of claims 2-7, 11-16, 20 and 21. Accordingly, Applicants respectfully request withdrawal of the rejection of claims 2-7, 11-16, 20 and 21 under 35 U.S.C. § 103(a).

V. Conclusion

It is respectfully urged that the subject application is patentable over Krause, Shah and Murayama and is now in condition for allowance. The Examiner is invited to call the undersigned at the below-listed telephone number if in the opinion of the Examiner such a telephone conference would expedite or aid the prosecution and examination of this application.

Respectfully submitted,

DATE: June 8, 2004



Stephen J. Walder, Jr.
Reg. No. 41,534
Yee & Associates, P.C.
P.O. Box 802333
Dallas, TX 75380
(972) 367-2001
Attorney for Applicants